




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 11/36	A1	(11) International Publication Number: WO 00/41078 (43) International Publication Date: 13 July 2000 (13.07.00)
(21) International Application Number: PCT/US99/31208 (22) International Filing Date: 30 December 1999 (30.12.99) (30) Priority Data: 09/224,921 4 January 1999 (04.01.99) US (71) Applicant: EMC CORPORATION [US/US]; 171 South Street, Hopkinton, MA 01748-9103 (US). (72) Inventor: SHAGAM, Eli; 1265 Beacon Street, Brookline, MA 02146 (US). (74) Agent: WALPERT, Gary, A.; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).		(81) Designated States: IL, JP, KR, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: ONLINE DEBUGGING AND TRACING SYSTEM AND METHOD

(57) Abstract

The invention relates to a method and apparatus for debugging software running in a target machine. A debugging set-up script is created in a host machine which defines trace point locations, and the variables to be returned to the host machine. The method sends the trace point locations and variables to the target machine where a stub program running in the target machine effects the modification of a software program in the target machine by inserting traps at the trace points. Data is collected using the stub program to ascertain variable values when a trace point is hit and the acquired variable data are stored in the target machine in a target machine buffer memory. The collected data is sent, at the request of the host machine, or at the end of a predetermined time, or when a pass-count is reached or at a time set by the target machine, to the host machine without stopping or interrupting operation of the target system.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

- 1 -

ONLINE DEBUGGING AND TRACING SYSTEM AND METHODBackground of the Invention

5 The invention relates generally to source code debugging methods and apparatus, and in particular, to the debugging of software programs using trace points set using a debugging tool.

 As computer programs become increasingly complex,
10 the programs will, more likely than not, contain errors or "bugs" that prevent the proper performance and operation of the program in its intended manner. The program is then debugged, which is the process of locating and correcting the errors in the program. In
15 complex programs, the debugging process be quite difficult and as a result one approach to debugging software uses the insertion of break points at locations within the code. The execution of the source code program then halts whenever a break point is encountered
20 to allow the programmer to observe the state of certain variables and, accordingly, the behavior of the program at the break point.

 In certain applications, stopping the program at a break point can be quite disruptive to the system. Thus,
25 for example, in a large shared data storage subsystem such as the EMC Symmetrix series of products, halting the program in effect stops the entire data reading, writing, and caching process. In this instance, not only can the flow of data be disrupted for on the order of several
30 minutes, or more, but the behavior of the system, and hence the analysis of an error, can be completely masked by the system stoppage because, for example, other external subsystems may then go into a recovery mode.

 It has been known to store data relating to trace
35 events in buffer memories within the Symmetrix device for later analysis. This approach has the advantage of not

- 2 -

interrupting the system while at the same time collecting data, as necessary, for later review and analysis. Such a system typically does not enable the user to dynamically alter or change the system software and trace
5 points, or to contemporaneously analyze the system software as a program continues.

As a result, in accordance with the present invention, a debugging system tool and method can be provided which enable immediate and certain interaction,
10 on line, with the software program operating, for example, in a large shared data storage system. In addition, the method and apparatus of the current invention advantageously provide online access to debug the system while, at the same time, not significantly
15 disrupting operation of the system so as to mask any error occurring in the operation of the source code.

Summary of the Invention

The invention relates to a method and apparatus for online debugging software running on a target
20 machine. The method features defining trace point locations (addresses) and data and variables to be collected at those trace points in the software at a host machine; sending the trace point locations and variables, for example in a debugging script, to the target machine,
25 maintaining a stub program in the target machine to perform the debugging script, collecting the data and variables at the predefined trace points, using the stub program, the data representing variables identified by the host machine debugging script when a trace point is
30 reached, and sending the collected data, online, at the request of the host machine, without significantly interrupting or stopping operation of the target machine. In a particular embodiment, the target machine is a disk drive controller.

- 3 -

In other aspects, the invention features determining, at the target machine, using an expression evaluation, locations for each variable for which data is to be collected when a trace point is reached and passed through, the variables being originally specified by the host machine debugging script in the form of numeric expressions from a compiler symbol table. In another aspect, the trace points may be automatically set at the host machine and the variables are automatically identified at the host machine.

In another aspect, in accordance with the invention, the collected data is directed to a buffer, in the target machine, for storage; and when the buffer overflows, the data is wrapped around the buffer thereby erasing the old data and replacing it with new data. In this way, only the last frames of data are kept in the buffer.

In another aspect, the apparatus of the invention provides for debugging software (the stub) running in a target machine and comprises a host machine, elements for defining, in the host machine, trace point locations and variables to be collected by the software, circuitry for sending the trace point locations and variables to the target machine, circuitry for running a stub program in the target machine, circuitry for collecting, using the stub program, data representing the variables selected by the user on the host machine. When a trace point is reached, the circuitry sends the collected data, online, to the host machine without stopping operation of the target machine.

In yet another aspect of the invention, the apparatus is a computer implemented apparatus for debugging, from a host computer, software running in a target machine. The invention provides for software media in both the target machine and the host computer to

- 4 -

implement the steps of defining, in the host computer, trace point locations and variables to be collected by the software on the target machine, sending the trace point locations and variables to the target machine, 5 running a stub program in the target machine, collecting, using stub program, data representing the variables selected by the user, for example, at the host computer when a trace point is reached and passed through, and providing the collected data, online, to the host 10 computer without stopping operation of the target machine.

The method and apparatus of the invention thus advantageously enable debugging of the target computer or disk controller with minimal disturbance to the operation 15 of the target machine.

Brief Description of the Drawings

Other features and advantages of the invention will be apparent to one practiced in this field from the following description of the invention taken together 20 with the drawings, in which:

Figure 1 is a general block diagram illustrating the system in which the invention has particular application;

Figure 2 is a block diagram illustrating in more 25 detail a typical environment in accordance with the invention;

Figure 3 is a diagrammatic block diagram illustrating a typical operation in accordance with the invention; and

30 Figure 4 is a flow chart illustrating operation of one embodiment in accordance with the invention.

Description of the Preferred Embodiments

Referring to Figure 1, the invention relates to a computer system wherein at least one, and more likely a

- 5 -

plurality of hosts 12a, 12b, ..., 12n, connect to a memory module system 14, such as the EMC Symmetrix disk array memory system. The memory module 14 acts as the interface between the host computers and a plurality of mass storage devices, such as, for example, disk drives 16a, 16b, ..., 16k. Data written by the host or read from the disk drive elements pass through the memory module system which acts as a two way communications path with substantial capabilities. For example, in some systems, the data from the host are uniformly striped across all or some of the disk storage devices; and in other systems, the data from the host are stored on the disk drives 16 according to a RAID protocol. In yet other embodiments of the invention, all of the data from a particular host can be stored on a single disk drive or in different logical volumes of the same or different disk drives, depending upon the nature and the source of the data and host. A host computer can also read data from one or more of the disk drive units.

When a problem arises which impairs performance of the system, for example, a non-recoverable software error, a decrease in throughput, or "bugs" in newly installed software, the problem can arise in either the host, the memory module, the disk drive elements, or in combinations thereof. In order to analyze and correct the problem, it is desirable not to bring down the customer's computer(s) or the controller, thereby placing them off-line and perhaps significantly impairing the customer's ability to do business. In accordance with the invention, therefore, when a software bug is suspected, for example, or to read system parameters at selected event times of the system operation, trace points are inserted in the software and data is collected without substantially imposing any significant performance degradation on the customer's system. As

- 6 -

described below, data can then be remotely collected for analysis without having to bring down, stop, or otherwise interfere substantially with the system operation.

Referring to Figure 2, in a particular embodiment according to the invention, the disk controller is configured to have a plurality of channel directors 30 (most often referred to a SCSI adapters when operating according to a SCSI protocol) connecting to a global memory 32 through which all data and commands flow. The global memory 32 is connected to a plurality of disk directors 34 (also typically SCSI adapters) which connect to the disk drives 16. In accordance with this particular embodiment of the invention, each channel director operates over channels 36 and 38 using a SCSI protocol. Each channel director 30 can be connected to one or more host computers over buses 36 (typically, one host I/O Controller per port 40). In the illustrated embodiment, it is the software operation which will be monitored and analyzed.

Referring now to Figure 3, in a diagrammatic representation, the host machine 400 is typically located in a site remote from the target machines 402. The host 400 and target 402 communicate over modems 404, 406 and a communications link 408. The host machine includes a source level debugger 410 which has access to a storage 412 containing source files 414, and the compiled symbol table 416 and binary executable file 418 for a program running in the target machine. Typically, the storage 412 contains many such collections of files for different versions of programs running on different target machines. The binary files, are identical to the binary executable program files at the target machine.

The source level debugger, using the symbol table 416, and under the control of either the user or an automatic trace point program, identifies the addresses

- 7 -

at which trace points are to be inserted into the executable program running at the target machine, as well as addresses for the variables (or numeric expressions to determine such addresses). That data is sent over
5 communications link 408 to the target machine. At the target machine the stub program uses that data and inserts traps at the trace point addresses, causing at each trace point, the necessary data to be collected and stored in a trace buffer 420. The stub program 422 is
10 included in the target machine as part of its operating programs. Thus, in general operation, a program is compiled and linked and loaded in binary executable form in the target machine while being stored in binary files 418 and symbol files 416. A user, or the system, then
15 defines the trace points and the data to be collected, and sends that set up information to the target machine. The appropriate data is then collected each time a trace point is hit and the trace program is terminated as described hereinafter. A "post mortum" analysis is then
20 performed on all or part of the data collected in the trace buffer.

Referring now to Figure 4, in operation, a source level debugging program is initiated at a host computer which is typically remote from the site of the target
25 machine. This is indicated at 98. The debugging program, under user control, opens a source code window, at the host computer, that is, a window on a display screen by which a debug information is created. This is indicated at step 100. The user then identifies a
30 program, running on the target machine, here controller system 14, and the debugger verifies that it has the same version of the program as that which is running on the target machine. This is indicated at 101. The user either automatically or manually sets a series of trace
35 points in the program to be debugged. This is indicated

- 8 -

at step 102. The manual setting of the trace points is typically performed at the source code level in the program. Along with each trace point the user either manually or automatically identifies the variables, data
5 for which are collected each time the trace point is reached or "hit". If the trace points are set automatically, a method such as that described in copending U.S. Patent Application No. 09/069,608, filed April 29, 1998, and entitled source code debugging tool
10 application, the contents of which are incorporated herein by reference, can be employed. Once the debug setup has been completed by the user, the host sends to the target machine the location of the trace points in the executable code existing on the target machine. The
15 format of the transmitted data allows numeric expressions to be sent from which the location of the desired variables can be derived at the target machine. This is indicated at 103. Typically the addresses and expressions are derived from the variables found in the
20 symbol table ordinarily created and stored by the compiler in creating the executable code of the program. The symbol table is available at the host machine. This address and expression information is transmitted at 104 to the target machine (as debugging set-up data) for use
25 in operation by the stub program there. The target and host systems can be connected, for example, over the internet, by modem, or a high speed communication bus. At the target machine, the so-called "stub" operates to implement the debugging set-up data at 105 and collects
30 the relevant data during operation of the software to be debugged. In this respect, the stub program inserts traps at the trace point addresses identified by the host system and collects the required data each time a "trap" is reached or hit.

- 9 -

As the program to be debugged at the target machine proceeds, the stub operates at each trace point address to trap the code and collect the required data. This is indicated at 110. When the trace point has been
5 hit, the stub operates to acquire the data from the target by evaluating the expressions provided by the host (the variable physical address can change from time to time), and store the acquired data in a target buffer. This is indicated at 112. If the target buffer fills, in
10 this illustrated embodiment of the invention, the buffer wraps around so that old and earlier collected data is overwritten.

The stub thus operates to collect the variable data specified in the debugging set-up even though it is
15 specified at the numeric or expression level. That is, the stub has an expression evaluator, working in reverse polish notation in the format of the preferred embodiment of the invention, which enables the stub to determine the address of the variable, or variables to be collected,
20 even though the address of the variables may change from time to time in the program.

Once all of the program variables have been collected, a pass-count for a trace point has been reached or at a time specified by the host or target
25 system has elapsed, or when a user generated access command is received by the target system, the collected and stored data can be off-loaded, in whole or in part as specified by the host machine, without substantially interrupting the operation of the target machine or the
30 monitored program, that is, while the target machine continues to operate. In a particular aspect, when the pass-count for a trace point is reached, the collection of data will automatically stop and the data will be made available to the host, either automatically or under user
35 control at the host. For example, for a pass count of

- 10 -

one, data will be collected when a particular trace point is first reached and can immediately be made available to the host for review and analysis. It should also be apparent that the collected data can be returned to the host while the additional data is being collected. This is indicated at step 120. The offloading process can be implemented and controlled by the host as indicated at step 130. In a particular embodiment, the host sends a search query to the target to obtain a limited, well defined, data download.

It is important to note that each time a trace point is reached, the program is trapped and the appropriate variables are collected and stored. This takes on the order of, for example, one millisecond. This is a significant improvement over, for example, those systems which, upon encountering a break point cause the program being monitored to stop, waiting for user input which can require interruptions of one, ten, or more minutes. As a result, the operation of the target machine much more closely resembles that of the machine without the trace points being implemented. Of course, the version of the software in the target machine must be identical to that known to the host computer.

Additions, subtractions, and other modifications of the disclosed preferred embodiment of the invention will be apparent to those practiced in the art and are within the skills of the following claims.

What is claimed is:

- 11 -

1. A method for debugging software running in a target machine comprising
 - defining, in a host machine, trace point locations and variables to be acquired by said software,
 - 5 sending the trace point locations and variables to the target machine,
 - running a stub program in the target machine,
 - collecting, using the stub program, data representing the variables selected by said host machine,
 - 10 when a trace point is reached, and
 - sending said collected data, online, to the host machine without stopping operation of said target machine.
2. The method of claim 1 further comprising
 - 15 determining at the target machine a location for each variable for which data is to be collected, for each trace point.
3. The method of claim 2 further comprising using the stub program at the target machine for collecting the
 - 20 identified data and delivering it to a storage buffer.
4. The method of claim 1 further comprising
 - automatically selecting said trace points at the host machine, and
 - automatically identifying the variables at the
 - 25 host machine.
5. The method of claim 1 further comprising
 - effecting traps in the software code at the target machine at said trace points.

- 12 -

6. The method of claim 1 further comprising
collecting data at said target machine until a
stop time set by one of the target and host machines.

7. The method of claim 1 further comprising
5 terminating said collecting of data when a trace point is
encountered a predetermined number of times.

8. The method of claim 6 further comprising
storing the collected data in a target machine
buffer, and
10 wrapping around said buffer if said target machine
buffer is filled.

9. A method for debugging software running in a
disk drive controller comprising
defining, in a host machine, trace point locations
15 and variables to be acquired by said software,
sending the trace point locations and variables to
the disk drive controller,
running a stub program in the disk drive
controller,
20 collecting, using the stub program, data
representing the variables selected by said host machine,
when a trace point is reached, and
providing said collected data, online, to the host
machine without stopping operation of said disk drive
25 controller.

10. The method of claim 9 further comprising
effecting traps in the software code at the disk drive
controller at said trace points while said disk drive
controller performs its normal read and write operations.

- 13 -

11. The method of claim 10 further wherein said providing step comprises providing said collected data to said host machine without interrupting operation of said disk drive controller.

5 12. The method of claim 9 further comprising terminating said collecting of data when a trace point is encountered a predetermined number of times.

13. Apparatus for debugging software running in a target machine comprising
10 means for defining, in a host machine, trace point locations and variables to be acquired by said software,
 means for sending the trace point locations and variables to the target machine,
 means for running a stub program in the target
15 machine,
 means for collecting, using the stub program, data representing the variables selected by said host machine, when a trace point is engaged, and
 means for sending said collected data, online, to
20 the host machine without stopping operation of said target machine.

14. The apparatus of claim 13 further comprising
 means for determining at the target machine a location for each variable for which data is to be
25 collected, at each trace point engagement.

15. The apparatus of claim 13 further comprising
 means for automatically selecting said trace points at the host machine, and
 means for automatically identifying the variables
30 at the host machine.

- 14 -

16. The apparatus of claim 13 further comprising means for effecting traps in the software code at the target machine at said trace points.

5 17. The apparatus of claim 13 further comprising means for collecting data at said target machine until a stop time set by one of the target and host machines.

18. The apparatus of claim 17 further comprising means for collecting data at said target machine
10 until a stop time set by one of the target and host machines.

19. The apparatus of claim 17 further comprising a target machine buffer,
means for storing the collected data in the target
15 machine buffer, and
means for wrapping around said buffer if said target machine buffer is filled.

20. A computer implemented apparatus for debugging, from a host computer, software running in a
20 target machine comprising software programs stored on magnetic media and requiring the steps of
defining in a host computer, trace point locations and variables to be acquired by said software,
sending the trace point locations and variables to
25 the target machine,
running a stub program in the target machine,
collecting, using the stub program, data representing the variables selected by said host computer, when a trace point is passed, and

- 15 -

sending said collected data, online, to the host computer without stopping operation of said target machine.

21. The apparatus of claim 19 further having
5 software comprising
determining, at the target machine, a location for each variable for which data is to be collected, at each trace point.

22. The apparatus of claim 20 further having
10 software comprising
using the stub program at the target machine for collecting the identified data and delivering it to a storage buffer.

23. The apparatus of claim 19 further having
15 software comprising
automatically selecting said trace points at the host computer, and
automatically identifying the variables at the host computer.

20 24. The apparatus of claim 19 further having
software comprising
effecting traps in the software codes at the target machine at said trace points.

25 25. The apparatus of claim 19 further having
software comprising
collecting data at said target machine until a stop time set by one of the target machine and host computers.

- 16 -

26. The apparatus of claim 24 further having
software comprising
storing the collected data in a target machine
buffer, and
5 wrapping around said buffer if said target machine
buffer is filled.

27. The method of claim 19 further comprising
terminating said collecting of data when a trace point is
encountered a predetermined number of times.

1/3




FIG. 1




FIG. 2

2/3




FIG. 3

3/3

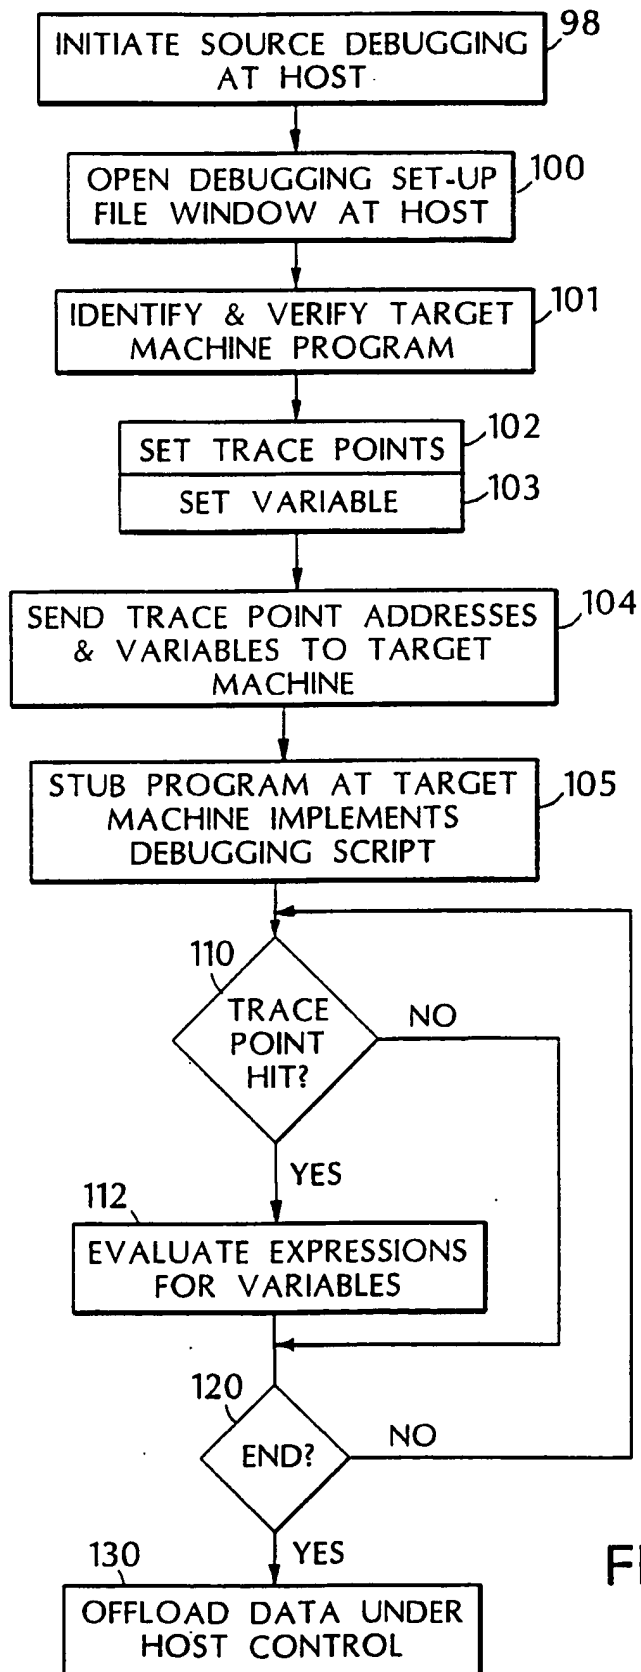


FIG. 4

INTERNATIONAL SEARCH REPORT

Int. National Application No.

PCT/US 99/31208

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F11/36

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	ANDREAS DANUSER ET AL.: "Was geschieht wann wo im Programm?" TECHNISCHE RUNDschau., vol. 82, no. 35, 31 August 1990 (1990-08-31), pages 48-57, XP000151488 HALLWAG VERLAG. BERN., CH ISSN: 1023-0823 page 50, right-hand column, line 20 -page 57, middle column, line 14 -----	1-7, 9-18, 20-25, 27



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

G document member of the same patent family

Date of the actual completion of the international search

17 May 2000

Date of mailing of the international search report

24/05/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Corremans, G